I Accidentally the Whole Internet

Ivana Bachmann
NICLabs, Universidad de Chile
Blanco Encalada 1975
Santiago de Chile
ivana@niclabs.cl

Fernando Morales
NICLabs, Universidad de Chile
Blanco Encalada 1975
Santiago de Chile
fernando@niclabs.cl

Alonso Silva
Bell Labs, Nokia
Centre de Villarceaux,
Route de Villejust,
91620 Nozay, France
alonso.silva@nokia.com

Javier Bustos-Jimenez
NICLabs, Universidad de Chile
Blanco Encalada 1975
Santiago de Chile
jbustos@niclabs.cl

arXiv:1605.08714v1 [cs.NI] 27 May 2016

ABSTRACT

Whether as telecommunications or power systems, networks are very important in everyday life. Maintaining these networks properly functional and connected, even under attacks or failures, is of special concern.

This topic has been previously studied with a whole network robustness perspective, modeling networks as undirected graphs (such as roads or simply cables). This perspective measures the average behavior of the network after its last node has failed.

In this article we propose two alternatives to well-known studies about the robustness of the backbone Internet: to use a supply network model and metrics for its representation (we called it the Go-Index), and to use robustness metrics that can be calculated while disconnections appear.

Our research question is: if a smart adversary has a limited number of strikes to attack the Internet, how much will the damage be after each one in terms of network disconnection?

Our findings suggest that in order to design robust networks it might be better to have a complete view of the robustness evolution of the network, from both the infrastructure and the users perspective.

Categories and Subject Descriptors

D.2.8 [**Software Engineering**]: Metrics—*complexity measures, performance measures*; E.1 [**Data**]: Data Structures—*graphs and networks*; G.2.2 [**Discrete Mathematics**]: Graph Theory—*network problems*

General Terms

Algorithms,Measurement,Reliability,Theory

Keywords

Complex Networks, Internet Backbone, Robustness Metrics

1. INTRODUCTION

Transportation, electrical and telecommunication networks, to name a few, have become fundamental for the proper functioning of the modern world. For that reason it has become of extreme importance that these systems remain operative. However these systems are prone to failure due to malfunctions, catastrophes or attacks.

All these structures can be studied through complex networks by representing the components of the structure by nodes and interactions among the components by edges.

Since their correct functioning requires that the network is properly connected it is of great importance to study their ability to resist failures (either unintentional or targeted attacks). This ability is called robustness.

In this work, we focus on the scenario of targeted attacks by an adversary. We notice that this scenario corresponds to an upper bound on the damage of (unintentional or intentional) failures.

We consider that an "adversary" should plan a greedy strategy aiming to maximize the damage with minimum number of strikes. Thus, in this article we discuss the performance of attacks based on the edge betweenness centrality metric [2] over the Internet Backbone (the network formed by Internet exchange points, IXP), and its correlation with what users perceive from such networks if they want to receive content from the major content provider (Google), with a metric called Go-Index.

We consider the Go-index measure with contains different supply network measures whose provider is Google. Just like economy uses the price of the Big Mac as a way of measuring purchasing power parity for its wide availability, here we measure the ability of the nodes to remain connected with Google for the same reason.

The idea to consider an IXP-based network as "the Internet backbone" is not new, previously has been used, for instance, as part of the "internet core" to study the inter-AS traffic patterns and an evolution of provider peering strategies [15], for optimize the content delivery from Google via direct paths [7] (in this case filtering IXP ASs because they facilitate connectivity between peers), and the Internet Backbone Market [4]. The importance of our study is based in with the use of the IXP network as a model for "backbone Internet", we can have a good approximation of Internet physical robustness. As far as the authors knowledge, this is the first time that the robustness of IXP network is studied.

The article is organized as follows, next section presents related work, followed by the methodology for building the IXP network, the attacking strategy using betweenness centrality and Go-Index (Sections 4 and 5). Conclusions are presented in Section 6.

2. RELATED WORK

To study the robustness of a network its evolution against failure must be analyzed. On real world situations networks may confront random failures as well as targeted attacks. For the latter, two main categories of attacking strategies have been defined: simultaneous and sequential attacks [11]. Simultaneous attacks chose a set of nodes and remove them all at once while sequential attacks chose a node to remove and given the impact of this removal it chooses another node, proceeding in turns.

Targeted attacks have been thoroughly studied to analyze network robustness. Holme *et al.* tested node degree and betweenness centrality strategies using simultaneous and sequential attacks. On [24] simultaneous attacks based on network centrality measures and random attacks were studied.

The stability of scale-free network under degree-based attacks was studied on [27]. Experimental results are shown in [13] who also consider sequential and simultaneous attacks as well as centrality measure strategies. To get closer to a real world strategy scenario on [10, 26] studied the resilience of scale-free networks to a variety of attacks with different amounts of information available to the attacker about the network.

On [3] the impact of the effectiveness of the attack under observation error was studied. More recently [25] studied sequential multi-strategy attacks using multiples robustness measures including the *Unique Robustness Measure* (R-index) [20].

The attacking strategies have been analyzed through the lens of an attacker, an adversary whose objective is to perform the most damage possible to the network. However the case of an adversary with a limited amount of strikes remain to be tested. Here this case is analyzed and an option to measure the robustness of a network in these circumstances is presented.

On [9] was found that targeted attack can be more effective when they are directed to bottlenecks rather than hubs. On [1] authors present partial values of R-index while nodes are disconnected, showing the importance of a well chosen robustness metric for performing the attacks.

The idea of planning a "network attack" using centrality measures has captured the attention of researchers and practitioners nowadays. For instance, Sterbenz et al. [22] used betweenness-centrality (*bcen*) for planning a network attack, calculating the *bcen* value for all nodes, ordering nodes from higher to lower *bcen*, and then attacking (discarding) those nodes in that order. They have shown that disconnecting only two of the top *bcen*-ranked nodes, their packet-delivery ratio is reduced to 60%, which corresponds to 20% more damage than other attacks such as random links or nodes disconnections, tracked by link-centrality and by node degrees.

A similar approach and results were presented by Çetinkaya et al. [6]. They show that after disconnecting only 10 nodes in a network of 100+ nodes the packet-delivery ratio is reduced to 0%. Another approach, presented as an improved network attack [18, 23], is to recalculate the betweenness-centrality after the removal of each node [12, 17]. They show a similar impact of non-recalculating strategies but discarding sometimes only half of the equivalent nodes.

In the study of resilience after edge removing, Rosenkratz et al. [19] study backup communication paths for network services defining that a network is "*k-edge-failure-resilient if no matter which subset of k or fewer edges fails, each resulting subnetwork is self-sufficient*" given that "*the edge resilience of a network is the largest integer k such that the network is k-edge-failure-resilient*".

For a better understanding of network attacks and strategies, see [12, 17, 18, 23].

3. BUILDING THE INTERNET BACKBONE GRAPH

Internet peering is the contract (formal or informal) between two autonomous systems (AS) that agree to exchange traffic (and traffic routes) through a physical link. In [8] authors present that "*The core of the Internet is a multi-tier hierarchy of Transit Providers (TPs). About 10-20 tier-1 TPs, present in many geographical regions, are connected with a clique of peering links. Regional (tier-2) ISPs are customers of tier-1 TPs. Residential and small business access (tier-3) providers are typically customers of tier-2 TPs*". Therefore, it is natural to think that the peering network is a coarse grained approximation of the Internet itself. Thus, we used it to model Internet for studying its robustness.

From `peeringdb.com` we collected the autonomous

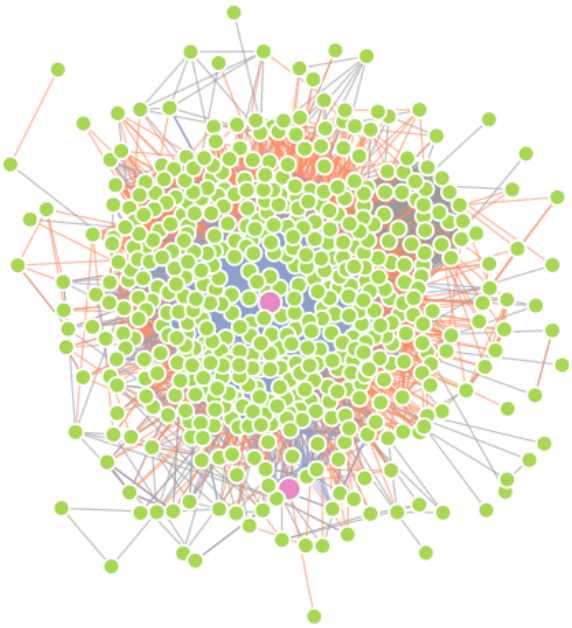

Figure 1: Peering Graph

systems (AS) from every Internet Exchange Point (IXP) and defined them as graph nodes. Therefore, an AS could belong to different IXPs and an IXP could have multiple ASs. Then, we connected the nodes if they fulfill at least one of the following rules:

- Physically linked ASs that exchange traffic
- ASs belonging to the same IXP
- ASs belonging to the same facility

Where we considered IXP as public peering and facility as private peering.

Figure 1 shows the resulting Graph, which has 522 nodes and 14,294 edges (orange edges are public peering, blue are private peering, and green are direct network connection). The resultant network has a well connected core network with some isolated nodes at the edges. In Figure 2 we present a degree distribution of nodes for our IPX Graph.

4. STUDYING THE ROBUSTNESS OF THE INTERNET BACKBONE

Betweenness centrality is a metric that determines the importance of an edge by looking at the shortest paths between all of the pairs of remaining nodes. Betweenness has been studied as a resilience metric for the routing layer [21] and also as a robustness metric for complex networks [13] and for internet autonomous systems networks [16] among others. Betweenness centrality has been widely studied and standardized as a

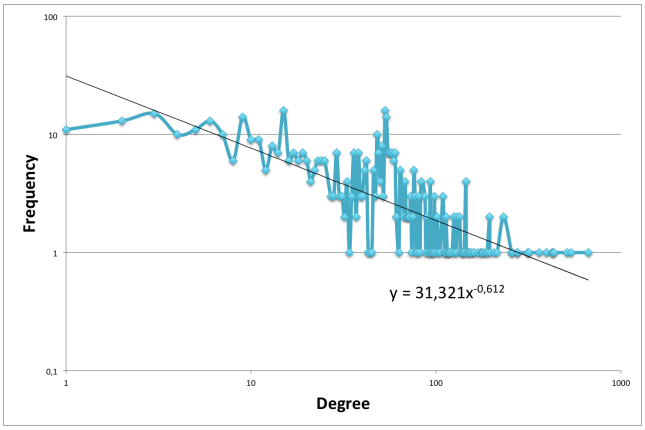

Figure 2: Degree distribution of IPX autonomous systems

comparison base for robustness metrics, thus in this study it will be used for performance comparison.

If we plan a network attack by disconnecting edges with a given strategy, it is widely accepted to compare it against the use of betweenness centrality metric, because the latter reflects the importance of an edge in the network [13]. These attack strategies are compared by means of the *Unique Robustness Measure* (R-index) [20], defined as:

$$R = \frac{1}{N}\sum_{Q=1}^{N} s(Q), \tag{1}$$

where N is the number of nodes in the network and $s(Q)$ is the fraction of nodes in the largest connected component after disconnecting edges using a given strategy. Therefore, the higher the R-index, the better in terms of robustness.

Instead of just comparing the robustness, after the removal of all of the edges, we would like to study the behavior of the attacks after only a few strikes. To do so, we define a variant of the R-index which takes into account only the first n strikes of an attack. Thus, for a simultaneous attack (where the nodes are ranked by a metric only once at the beginning), the R_n-index is defined as:

$$R_n = \frac{1}{n}\sum_{Q=1}^{n} s(Q). \tag{2}$$

That is, the area under the curve produced by the largest connected component ratio (compared to the whole network) until n.

For a sequential attack, the order of node disconnection is recomputed after each disconnection. Similar to the R-index, notice that the lower the R_n-index, the more effective the attack is, since that gives us a higher reduction of robustness.

Results are shown in Figure 3. We tested sequential

attacks: At each strike, the next edge to disconnect was the one with the highest betweenness value. The figure shows the behavior of the R_n-index in our IPX Graph. The strategy proves to be very effective in attacks, disconnecting half of the network removing only 20% of the edges, more than 30% of the nodes after removing 10% of edges, and 10% of nodes after 1% of edges.

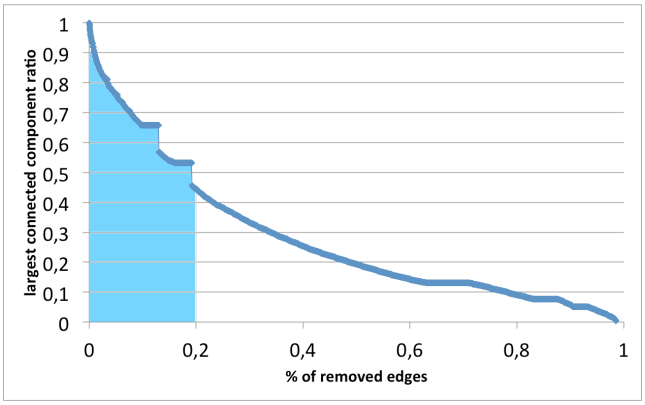

Figure 3: % of the largest component size compared to the original network. In the plot, $R_{20\%}$-index is marked in cyan.

5. THE GO-INDEX: IF YOU CANNOT SEE GOOGLE, YOU ARE NOT CONNECTED

Google has been reported as having the 40% of Internet traffic in 2013[1]. Thus, given the Google peering policies[2] and knowing the Google policies to interconnect their datacenters [14], we can also study Internet as an information supply network, adapting to Internet the metrics presented in [28]:

1. **Supply Availability** (SAR): The percentage of ASs that have access to Google from at least one of its ASs.

2. **Network Connectivity** (NetCON): The number of ASs in the largest functional sub-network, in which there is a path between any pair of ASs and there exist at least one of the Google ASs.

3. **Best Delivery Efficiency** (BDE): The reciprocal of the average of each demand AS's shortest supply path length to its nearest Google AS. Values go from 1 (everyone is connected directly with a Google AS) to 0 (there are only Google ASs in the network).

4. **Average Delivery Efficiency** (ADE): The average inverse supply path length for all possible {Network AS,Google AS} pairs, adjusted by a weighting factor for each path (in our study all Google ASs have the same importance). In this case, values go from 2 (everyone is connected to both Google ASs directly) to 0 (nobody is connected with Google ASs).

Notice that Google delegated some services at ISPs autonomous systems [5], nevertheless they must eventually connect with Google backbone for updating. We called the tuple {1,2,3,4} the *Go-Index*, that is, the supply network measures whose provider is Google.

Using the same attack strategy from previous section, we calculated the Go-Index after edge removal (removing the edge with higher betweenness). The results are presented in Figure 4 (Supply Availability), 5 (Network Connectivity), 6 (Best Delivery Efficiency), and 7 (Average Delivery Efficiency).

The first two metrics are very related with the largest connected component ratio, which in this study include at least one of the two Google ASs, that is, AS15169 and AS36040 (marked in pink at Figure 1, the former in the center and the latter in the edge of the network). Therefore, no new information are provided by those metrics.

The following two metrics are based in "how far is Google from a given autonomous system", but BDE considers only the connected component that includes Google ASs, by itself it has no information about the isolated portion of the network that cannot reach the Google ASs, therefore BDE improves when large subnets are disconnected from the core network that contains the Google ASs.

Note that for a user inside that core network the main content provider always exists and there are no indications that the network is falling apart (or losing half of its members, as produced by eliminating 20% of its edges), this can be appreciated through BDE. Nevertheless, the big picture is different: after having only the 5% of the network disconnected one of Google ASs is isolated, showing that from this point the supply network is only maintained by AS15169. The perception error is corrected when ADE is used because it includes all nodes in its calculus.

Then, the Go-Index accomplish its objectives, reflecting both infrastructure (SAR+NetCON+ADE) and user perception (BDE+ADE), for Internet robustness studies.

6. CONCLUSIONS AND FUTURE WORK

In this paper we have presented how robust the Internet backbone (the peering AS network) would be if an adversary can choose wisely which physical link he will cut (or if a very unlucky accident happen). Following the recommendations, the chosen one would be the edge with higher betweenness centrality value.

[1]See the Forbes article at `http://goo.gl/aHdeiN`
[2]See `https://peering.google.com/#/options/peering`.

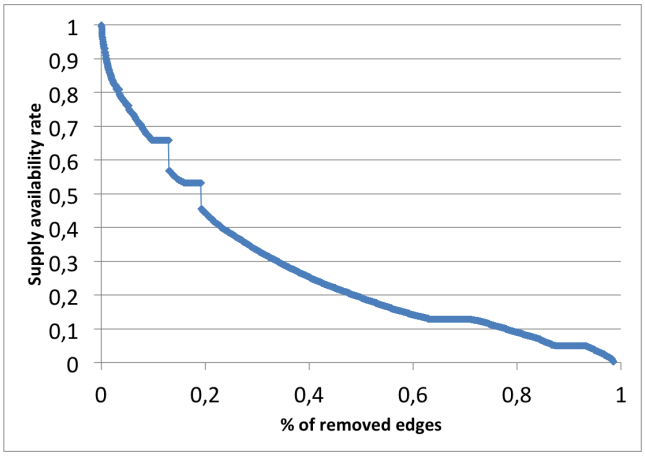

Figure 4: Supply Availability Ratio

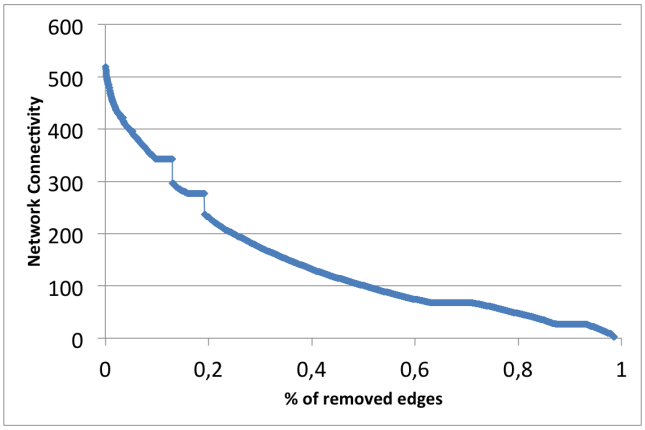

Figure 5: Network Connectivity

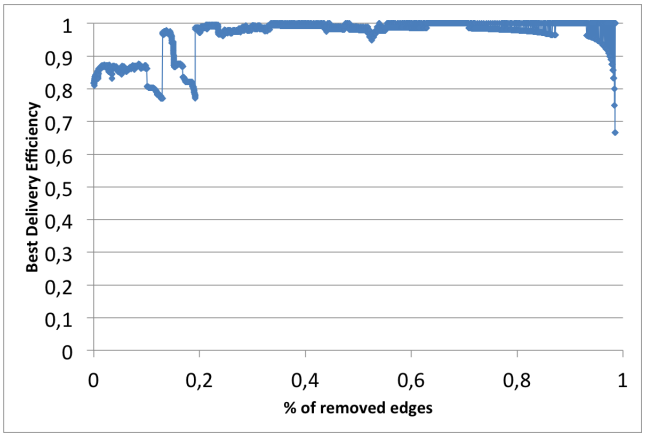

Figure 6: Best Delivery Efficiency

Using this strategy the adversary is capable of disconnecting half of the network by removing only 20% of the edges, more than 30% of the nodes after removing 10% of the edges, and 10% of nodes after 1% of the edges as we have seen with the values of R_n index in Section 4.

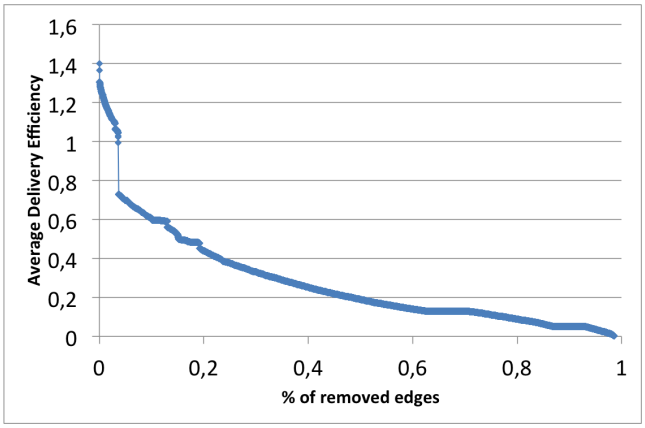

Figure 7: Average Delivery Efficiency

Furthermore, we consider the Internet as a (information) supply network, and considering that Google is the main Internet content provider, and propose to study the Internet backbone with the Go-Index (the adaptation of metrics presented in [28] for supply networks).

The metrics used were Supply Availability, Network Connectivity (both highly correlated with the larger connected component ratio), Best Delivery Efficiency, and Average Delivery Efficiency.

If only Best Delivery Efficiency is considered, the network can be declared robust because a user located inside the core network, which is always connected with a Google AS, will not perceive that the network is being disconnected. The Go-index will correct that perception since it also contains the Average Delivery Efficiency which includes all nodes in its calculus.

As future work we plan to apply similar studies to other Internet infrastructures, such as country-based fiber interconnection, submarine Internet cables, etc. Also, we plan to improve the metrics for robustness reflecting both the infrastructure part (such as Rn index) and the user perception (such as Best Delivery Efficiency/Average Delivery Efficiency).

7. REFERENCES

[1] I. Bachmann, P. Reyes, A. Silva, and J. Bustos-Jimenez. Miuz: measuring the impact of disconnecting a node. In *2015 34th International Conference of the Chilean Computer Science Society (SCCC)*, pages 1–6, Nov 2015.

[2] N. I. Bersano-Méndez, S. E. Schaeffer, and J. Bustos-Jiménez. Metrics and models for social networks. In *Computational Social Networks*, pages 115–142. Springer, 2012.

[3] L. B. Booker. The effects of observation errors on the attack vulnerability of complex networks. Technical report, DTIC Document, 2012.

[4] P. Buccirossi, L. Ferrari Bravo, and P. Siciliani. Competition in the internet backbone market. *World Competition*, 28(2):233–252, 2005.

[5] M. Calder, X. Fan, Z. Hu, E. Katz-Bassett, J. Heidemann, and R. Govindan. Mapping the expansion of google's serving infrastructure. In *Proceedings of Internet Measurement Conference*, pages 313–326. ACM, 2013.

[6] E. K. Çetinkaya, D. Broyles, A. Dandekar, S. Srinivasan, and J. P. Sterbenz. Modelling communication network challenges for future internet resilience, survivability, and disruption tolerance: A simulation-based approach. *Telecommunication Systems*, 52(2):751–766, 2013.

[7] Y.-C. Chiu, B. Schlinker, A. B. Radhakrishnan, E. Katz-Bassett, and R. Govindan. Are we one hop away from a better internet? In *Proceedings of Internet Measurement Conference*, pages 523–529. ACM, 2015.

[8] A. Dhamdhere and C. Dovrolis. The internet is flat: Modeling the transition from a transit hierarchy to a peering mesh. In *Proceedings of Co-NEXT*, pages 21:1–21:12, New York, NY, USA, 2010. ACM.

[9] E. Estrada. Network robustness to targeted attacks. the interplay of expansibility and degree distribution. *The European Physical Journal B-Condensed Matter and Complex Systems*, 52(4):563–574, 2006.

[10] L. K. Gallos, R. Cohen, F. Liljeros, P. Argyrakis, A. Bunde, and S. Havlin. Attack strategies on complex networks. In *Computational Science*, pages 1048–1055. Springer, 2006.

[11] P. Holme, B. J. Kim, C. N. Yoon, and S. K. Han. Attack vulnerability of complex networks. *Physical Review E*, 65(5):056109, 2002.

[12] P. Holme, B. J. Kim, C. N. Yoon, and S. K. Han. Attack vulnerability of complex networks. *Physical Review E*, 65(5):056109, 2002.

[13] S. Iyer, T. Killingback, B. Sundaram, and Z. Wang. Attack robustness and centrality of complex networks. *PloS one*, 8(4):e59613, 2013.

[14] S. Jain, A. Kumar, S. Mandal, J. Ong, L. Poutievski, A. Singh, S. Venkata, J. Wanderer, J. Zhou, M. Zhu, et al. B4: Experience with a globally-deployed software defined wan. *ACM SIGCOMM Computer Communication Review*, 43(4):3–14, 2013.

[15] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian. Internet inter-domain traffic. *ACM SIGCOMM Computer Communication Review*, 41(4):75–86, 2011.

[16] P. Mahadevan, D. Krioukov, M. Fomenkov, X. Dimitropoulos, A. Vahdat, et al. The internet as-level topology: three data sources and one definitive metric. *ACM SIGCOMM Computer Communication Review*, 36(1):17–26, 2006.

[17] W. Molisz and J. Rak. End-to-end service survivability under attacks on networks. *Journal of Telecommunications and Information Technology*, pages 19–26, 2006.

[18] J. Rak and K. Walkowiak. Survivability of anycast and unicast flows under attacks on networks. In *International Congress on Ultra Modern Telecommunications and Control Systems*, pages 497–503. IEEE, 2010.

[19] D. J. Rosenkrantz, S. Goel, S. Ravi, and J. Gangolly. Resilience metrics for service-oriented networks: A service allocation approach. *Services Computing, IEEE Transactions on*, 2(3):183–196, 2009.

[20] C. M. Schneider, A. A. Moreira, J. S. Andrade, S. Havlin, and H. J. Herrmann. Mitigation of malicious attacks on networks. *Proceedings of the National Academy of Sciences*, 108(10):3838–3841, 2011.

[21] P. Smith, D. Hutchison, J. P. Sterbenz, M. Scholler, A. Fessi, M. Karaliopoulos, C. Lac, and B. Plattner. Network resilience: a systematic approach. *Communications Magazine, IEEE*, 49(7):88–97, 2011.

[22] J. P. Sterbenz, E. K. Cetinkaya, M. Hameed, A. Jabbar, J. P. Rohrer, et al. Modelling and analysis of network resilience. In *Third International Conference on Communication Systems and Networks*, pages 1–10. IEEE, 2011.

[23] A. Sydney, C. Scoglio, M. Youssef, and P. Schumm. Characterising the robustness of complex networks. *International Journal of Internet Technology and Secured Transactions*, 2(3-4):291–320, 2010.

[24] S. Trajanovski, J. Martín-Hernández, W. Winterbach, and P. Van Mieghem. Robustness envelopes of networks. *Journal of Complex Networks*, 1(1):44–62, 2013.

[25] M. Ventresca and D. Aleman. Network robustness versus multi-strategy sequential attack. *Journal of Complex Networks*, 3(1):126–146, 2015.

[26] J. Wu, H.-Z. Deng, Y.-J. Tan, and D.-Z. Zhu. Vulnerability of complex networks under intentional attack with incomplete information. *Journal of Physics A: Mathematical and Theoretical*, 40(11):2665, 2007.

[27] A. Yehezkel and R. Cohen. Degree-based attacks and defense strategies in complex networks. *Physical Review E*, 86(6):066114, 2012.

[28] K. Zhao, A. Kumar, and J. Yen. Achieving high robustness in supply distribution networks by rewiring. *Engineering Management, IEEE Transactions on*, 58(2):347–362, 2011.